

Efficient quantization for average consensus

Ruggero Carli, Fabio Fagnani, Paolo Frasca, Sandro Zampieri

March 7, 2009

Abstract

This paper presents an algorithm which solves exponentially fast the average consensus problem on strongly connected network of digital links. The algorithm is based on an efficient *zooming-in/zooming-out* quantization scheme.

1 Introduction

Problems of coordination, and in particular of distributed averaging, have become an important subject of research for control theorists [16]. In the average consensus problem, N agents are given one number each, and have to compute their average. The interest of the problem comes from the communication limitations the agents are subject to, which require the design of an iterative algorithm. Usually, the agents are allowed to communicate with a restricted number of neighbors, and this is described by a suitable directed graph. In this work, we present an average consensus algorithm in which the agents can communicate with their neighbors only via rate constrained digital channels. We propose an efficient encoding/decoding strategy, based on the so called *zooming-in/zooming-out* quantization method [3, 17, 14], and we give theoretical and simulation results on its performance. The algorithm solves the average consensus problem with arbitrary precision in exponential time, using a quantizer with number of states which is finite, but can depend on N .

1.1 Related works

The constraint of quantization of messages, or of states, has been considered in several recent papers [18, 9, 5, 11, 13, 12, 15, 1]. However, existing algorithms in the literature are not able to achieve average consensus with arbitrary precision, with the notable exception of [4], which uses logarithmic quantizers, together with a coding scheme similar to the one in the present paper. The zooming-in/zooming-out algorithm and the main convergence result of this paper appeared, in slightly different version, in [6] and in the thesis [10]. Note that recently [2] a modification of the present algorithm has been proved to converge under a different condition.

1.2 Outline

The problem of average consensus is formally presented in Section 2. The algorithm, and the main result, are in Section 4. In Section 5 we present several simulations, which show an encouraging performance, beyond the scope of the proven result. Finally, Section 6 points out some potential developments.

2 Average consensus

Let there be a directed graph $\mathcal{G} = (V, E)$, with $V = \{1, \dots, N\}$ a set of dynamical systems, that we call agents. We assume that only if $(j, i) \in E$, j can transmit information about its state to i . Let us consider, for $t \in \mathbb{Z}_{\geq 0}$, the following discrete-time state equations

$$x_i(t+1) = x_i(t) + u_i(t) \quad i = 1, \dots, N$$

where $x_i(t) \in \mathbb{R}$ is the state of the i -th system/agent, and $u_i(t) \in \mathbb{R}$ is the control input. More compactly we can write

$$x(t+1) = x(t) + u(t) \tag{1}$$

where $x(t), u(t) \in \mathbb{R}^N$.

Definition 1 (Average consensus problem) *Given the system (1), the average consensus problem consists in designing a sequence of control inputs $u(t)$ yielding the consensus of the states, namely*

$$\lim_{t \rightarrow +\infty} x(t) = x_{\text{ave}}(0)\mathbf{1}, \tag{2}$$

where $\mathbf{1}$ is a N -vector whose entries are 1.

The important case we restrict to, and which has been mostly studied in literature, is a static linear feedback

$$u(t) = Kx(t) \quad K \in \mathbb{R}^{N \times N} \tag{3}$$

In such case the system (1) is given by the following closed loop system

$$x(t+1) = (I + K)x(t). \tag{4}$$

The matrix $I + K$ is commonly called a Perron matrix, and denoted by P . The problem of designing the controller thus reduces to find a matrix K adapted to \mathcal{G} , in the sense that if $(j, i) \notin E$, then $K_{ij} = 0$. Methods for design are known in the literature [8]. We shall only recall here that the average consensus problem is solved if and only if the matrix P is a doubly stochastic matrix.

2.1 Quantization

We can make the framework (3) more general if we assume that the agents can not directly access their neighbors' states, but can obtain estimates or approximations of them. Let $\hat{x}_j^i(t)$ be the estimate agent i has of the state of agent j , at time t . The control input $u_i(t)$ assumes the following form

$$u_i(t) = \sum_{j=1}^N K_{ij} \hat{x}_j^i(t). \quad (5)$$

Due to the possibility that $\hat{x}_j^i(t) \neq x_j(t)$, such a control is not in general guaranteed to yield convergence to consensus, nor to preserve the average of states.

Most of the literature on the consensus problem assumes that the communication channels between the nodes allows to communicate real numbers with no errors. In practical applications this can be an unrealistic assumption: the communication is in most cases digital and the channel is noisy. There can be strict bandwidth constraints, as well as communication delays, and packet losses due to interferences and erasures. These limitations are mostly likely to be significant for a network communicating in a wireless fashion, which is nowadays a typical choice. In this paper, we concentrate on modeling the links as instantaneous and lossless digital channels. This choice forces a quantization on the real numbers that agents have to transmit, but keeps aside the issues of delays and packet losses.

3 Encoding-decoding schemes

From now on, let us assume to have a strongly connected directed graph \mathcal{G} and an average consensus controller K adapted to \mathcal{G} . In the extended communication framework (5), the control input $u_i(t)$ of agent i has the following form

$$u_i(t) = \sum_{j=1}^N K_{ij} \hat{x}_j^i(t), \quad (6)$$

where \hat{x}_j^i is the estimate of the state x_j which has been built by the agent i . The point now is to describe how the estimates are built. In other works, [11], [7], we have assumed that the estimate be just the transmitted message. Now, instead, we look for a non trivial way of constructing estimates.

Suppose that the j -th agent sends to the i -th agent, through a digital channel, at each time instant t , a symbol $s_{ij}(t)$ belonging to a finite or countable alphabet \mathcal{S}_{ij} , called the *transmission alphabet*. It is assumed that the channel is reliable, that is each symbol transmitted is received without error. In general, the structure of the coder by which the j -th agent produces the symbol to be sent to the i -th agent can be described by the following equations

$$\begin{cases} \xi_{ij}(t+1) &= F_{ij}(\xi_{ij}(t), s_{ij}(t)) \\ s_{ij}(t) &= Q_{ij}(\xi_{ij}(t), x_j(t)) \end{cases} \quad (7)$$

where $s_{ij}(t) \in \mathcal{S}_{ij}$, $\xi_{ij}(t) \in \Xi_{ij}$, $Q_{ij} : \Xi_{ij} \times \mathbb{R} \rightarrow \mathcal{S}_{ij}$, and $F_{ij} : \Xi \times \mathcal{S}_{ij} \rightarrow \Xi_{ij}$. The decoder, run by agent i , is the system

$$\begin{cases} \xi_{ij}(t+1) &= F_{ij}(\xi_{ij}(t), s_{ij}(t)) \\ \hat{x}_j^i(t) &= H_{ij}(\xi_{ij}(t), s_{ij}(t)), \end{cases} \quad (8)$$

where $H_{ij} : \Xi_{ij} \times \mathcal{S}_{ij} \rightarrow \mathbb{R}$. The set Ξ_{ij} serves as *state space* for the coder/decoder, whereas the maps F_{ij} , Q_{ij} , H_{ij} represent, respectively, the *coder/decoder dynamics*, the *quantizer function*, and the *decoder function*. Coder and decoder are jointly initialized at $\xi_{ij}(0) = \xi_0$.

In general, one may have different encoders at agent j , according to the various agents the agent j wants to send its data. For the sake of simplicity, we assume in the sequel that each system uses the same encoder for all data transmissions. Thus, each agent j broadcasts the same symbol $s_j(t)$ to all its neighbors. In this case every agent receiving data from j obtains the same estimate of $x_j(t)$, which we shall denote by $\hat{x}_j(t)$. In this way, by letting $F_j = F_{ij}$, $H_j = H_{ij}$, $Q_j = Q_{ij}$ and $\Xi_j = \Xi_{ij}$, the previous coder/decoder couple can be represented by the following *state estimator with memory*

$$\begin{cases} \xi_j(t+1) &= F_j(\xi_j(t), s_j(t)) \\ s_j(t) &= Q_j(\xi_j(t), x_j(t)) \\ \hat{x}_j(t) &= H_j(\xi_j(t), s_j(t)) \end{cases} \quad (9)$$

Moreover (6) assumes the following form

$$u_i(t) = \sum_{j=1}^N K_{ij} \hat{x}_j(t), \quad (10)$$

and the system evolution is

$$x(t+1) = x(t) + K \hat{x}(t). \quad (11)$$

Remark 1 (Need for synchrony) Remark that, from the point of view of the implementation, both agent j and its neighbors have to run identical copies of the system $\xi_j(t)$. This clearly implies an intrinsic weakness of the scheme with respect to failures: if messages are lost, the synchrony of these copies is lost as well, and the algorithm does not work.

In change of this potential drawback, the scheme has very interesting convergence properties.

4 Zooming-in zooming-out strategy

Our strategy is inspired by the quantized stabilization technique proposed in [3], which is called zooming-in/zooming-out strategy. In this case the information exchanged between the agents is quantized by scalar uniform quantizers which

assume values in a finite set, and can be described as follows. For $m \in \mathbb{N}$ define the set of quantization levels

$$\mathcal{S}_m = \{-1 + \frac{2\ell - 1}{m} \mid \ell \in \{1, \dots, m\}\} \cup \{-1, 1\}.$$

The corresponding *uniform quantizer* $q^{(m)} : \mathbb{R} \rightarrow \mathcal{S}_m$ is as follows. Let $x \in \mathbb{R}$ then

$$q^{(m)}(x) = -1 + \frac{2\ell - 1}{m}$$

if $\ell \in \{1, \dots, m\}$ satisfies $-1 + \frac{2(\ell-1)}{m} \leq x \leq -1 + \frac{2\ell}{m}$, otherwise $q^{(m)}(x) = 1$ if $x > 1$ and $q^{(m)}(x) = -1$ if $x < -1$.

It is easy to see that the quantizer $q^{(m)}$ enjoys the following property.

Lemma 1 (Quantization error) *Given $z \in \mathbb{R}$ and $l \in \mathbb{R}_{>0}$ such that $|z| \leq l$, it holds*

$$|z - lq^{(m)}\left(\frac{z}{l}\right)| \leq \frac{l}{m}. \quad (12)$$

Let $m \in \mathbb{N}$, $k_{\text{in}} \in]0, 1[$, and $k_{\text{out}} \in]1, +\infty[$. The *zooming-in/zooming-out* encoder/decoder, with parameters $m, k_{\text{in}}, k_{\text{out}}$, is defined by the alphabet $\mathcal{S} = \mathcal{S}_m$, the state space $\Xi_j = \mathbb{R} \times \mathbb{R}_{>0}$, such that $\xi_j(t) = (\hat{x}_j(t), l_j(t))$, and the dynamics

$$\begin{aligned} \hat{x}_j(0) &= 0 \\ \hat{x}_j(t+1) &= \hat{x}_j(t) + l_j(t+1)s_j(t+1) \quad \forall t \geq 0, \end{aligned}$$

and

$$\begin{aligned} l_j(0) &= l_0 \in \mathbb{R} \\ l_j(1) &= l_0 \\ l_j(t+1) &= \begin{cases} k_{\text{in}}l_j(t) & \text{if } |s_j(t)| < 1 \\ k_{\text{out}}l_j(t) & \text{if } |s_j(t)| = 1 \end{cases} \quad \forall t > 0. \end{aligned}$$

The sent symbol is

$$s_j(t) = q^{(m)}\left(\frac{x_j(t) - \hat{x}_j(t-1)}{l_j(t)}\right) \quad \forall t > 0.$$

Let us comment on such definition. Remark that the first component of the coder/decoder state contains $\hat{x}(t)$, the estimate of $x(t)$. The transmitted messages contain a quantized version of the estimation error $x_j(t) - \hat{x}_j(t-1)$ scaled by the factor $l_j(t)$. Accordingly, the second component of the coder/decoder state, l_j , is referred to as the *scaling factor*: it increases when $|x_j(t) - \hat{x}_j(t-1)| > l_j(t)$ (“zooming out step”) and decreases when $|x_j(t) - \hat{x}_j(t-1)| \leq l_j(t)$ (“zooming in step”).

We can now prove the main theorem of this chapter.

Theorem 2 (Convergence) Consider the system (11) with the zooming coding. Let ρ be the essential spectral radius of $I + K$. Suppose that $\rho < k_{\text{in}} < 1$, $m \geq \frac{(4+3k_{\text{in}})\sqrt{N}}{k_{\text{in}}(k_{\text{in}}-\rho)}$ and that $l_0 > \frac{2(\rho+2)\|x(0)\|}{k_{\text{in}} - \frac{3\sqrt{N}}{m}}$. Then, for any initial condition $x(0) \in \mathbb{R}^N$,

$$\lim_{t \rightarrow +\infty} x(t) = \lim_{t \rightarrow +\infty} \hat{x}(t) = x_{\text{ave}} \mathbf{1},$$

and the convergence is exponential, with rate not bigger than k_{in} .

Proof: First, observe that $\mathbf{1}^* K = 0$, and then the states update rule preserves the average of the state $x(t)$ at each time step. Hence if we prove convergence to a consensus value, that will be the average. Let us rewrite the overall system as the coupling of the three dynamical systems on \mathbb{R}^N ,

$$\begin{aligned} l_j(0) &= l_0 & \forall j \in V \\ l_j(1) &= l_0 & \forall j \in V \\ l_j(t+1) &= \begin{cases} k_{\text{in}} l_j(t) & \text{if } |x_j(t+1) - \hat{x}_j(t)| < 1 \\ k_{\text{out}} l_j(t) & \text{if } |x_j(t+1) - \hat{x}_j(t)| \geq 1 \end{cases} & \forall j \in V, t \in \mathbb{N} \end{aligned}$$

$$\begin{aligned} x(0) &= x_0 \\ x(t+1) &= x(t) + K \hat{x}(t) & \forall t \in \mathbb{Z}_{\geq 0} \end{aligned}$$

$$\begin{aligned} \hat{x}_j(0) &= 0 & \forall j \in V \\ \hat{x}_j(t+1) &= \hat{x}_j(t) + l_j(t+1) q^{(m)} \left(\frac{x_j(t+1) - \hat{x}_j(t)}{l_j(t+1)} \right) & \forall j \in V, t \in \mathbb{Z}_{\geq 0}. \end{aligned}$$

Their coupling implies that the following facts are equivalent:

- (i) $\lim_{t \rightarrow +\infty} l(t) = 0$;
- (ii) $\lim_{t \rightarrow +\infty} x(t) = x_{\text{ave}} \mathbf{1}$;
- (iii) $\lim_{t \rightarrow +\infty} (\hat{x}(t) - x(t)) = 0$.

The idea of the proof is the following: we show that under the assumptions there always happen zooming in steps. This implies $l(t) = k_{\text{in}}^{t-1} l_0 \mathbf{1}$ and implies in turn convergence to average consensus.

By definition, there are only zooming in steps if

$$|x_j(t+1) - \hat{x}_j(t)| < l_j(t+1) \quad \forall i \in V, t \in \mathbb{Z}_{\geq 0} \quad (13)$$

or equivalently if

$$|x_j(t+1) - \hat{x}_j(t)| < k_{\text{in}}^t l_0 \quad \forall i \in V, t \in \mathbb{Z}_{\geq 0}. \quad (14)$$

Let us define $\tilde{y}(t) = Kx(t)$ and $e(t) = x(t) - \hat{x}(t)$, and remark that (14) is equivalent to

$$|\tilde{y}(t) - (K - I)e(t)| < k_{\text{in}}^t l_0, \quad (15)$$

where the inequality is meant componentwise.

To prove this fact, we proceed by strong induction over t , proving that for all $t \in \mathbb{Z}_{\geq 0}$

$$\|\tilde{y}(t)\| \leq k_{\text{in}}^t l_0 \quad (16)$$

$$\|\tilde{y}(t) - (K - I)e(t)\| \leq k_{\text{in}}^t l_0. \quad (17)$$

Since $\|\tilde{y}(t) - (K - I)e(t)\|_{\infty} \leq \|\tilde{y}(t) - (K - I)e(t)\|$, the inequality (17) implies (14) and then convergence.

Let us thus check the validity of Equations (16) and (17) for $t = 0$, $t = 1$, and for a general $t + 1$, with $t \in \mathbb{N}$, given it true for all $t' < t$. We shall use the recursion

$$\tilde{y}(t + 1) = (I + K)\tilde{y}(t) - K^2 e(t),$$

as well as Lemma 1, which implies $\|e(t)\| \leq \frac{\sqrt{N}}{m} l_0$.

Using triangle and submultiplicative inequality of norms, and $\|I + K\| = \rho$, $\|K\| \leq 2$, and $\|K - I\| \leq 3$, we get that

$t = 0$. (16) holds if $\|Kx_0\| < l_0$, and (17) holds if $\|x_0\|_{\infty} < l_0$.

$t = 1$. (17) holds if

$$2(2 + \rho)\|x_0\| + 3\frac{\sqrt{N}}{m}l_0 < k_{\text{in}}l_0, \quad (18)$$

and (16) holds if

$$\|Kx_0\| < k_{\text{in}}l_0. \quad (19)$$

$t > 1$. Assuming (16) and (17) true for time steps t and $t - 1$, we obtain that (17) holds for the time step $t + 1$ if

$$\rho k_{\text{in}}^t + 4\frac{\sqrt{N}}{m}k_{\text{in}}^{t-1}l_0 + 3\frac{\sqrt{N}}{m}k_{\text{in}}^t l_0 < k_{\text{in}}^{t+1}l_0,$$

that is if

$$\rho k_{\text{in}} + 4\frac{\sqrt{N}}{m} + 3\frac{\sqrt{N}}{m}k_{\text{in}} < k_{\text{in}}^2. \quad (20)$$

Instead, (16) holds if

$$\rho k_{\text{in}} + 4\frac{\sqrt{N}}{m} < k_{\text{in}}^2. \quad (21)$$

The system of conditions (18)-(19)-(20)-(21) is satisfied when the assumptions of the theorem hold. \square

This result is quite conservative in proving convergence, since it guarantees that since the beginning only zooming in steps happen. This is clearly restrictive, since there is no need for $l(t)$ to be monotonic to converge to zero.

5 Simulation results

In this section we show by simulations the properties of the algorithm in terms of (speed of) convergence, as it depends on the parameters of the method, $k_{\text{in}}, k_{\text{out}}$ and m , and on the topology of the graph. Indeed, the convergence result we obtained, Theorem 2, gives sufficient conditions on the parameters, depending on ρ , and on N , and thus on the graph. Convergence is possible, provided the conditions, at exponential speed, with rate $k_{\text{in}} > \rho$. Simulations are worth of interest, since they demonstrate that convergence is possible also *outside the scope of Theorem 2*, and with a *speed which can be faster than the linear algorithm with ideal communication* (4).

Roughly speaking, we can identify two different regimes for the evolution of the algorithm, depending on the zooming-in rate k_{in} which is enforced. If $k_{\text{in}} > \rho$, the zooming-in follows the natural contraction rate of the system, and according to the spirit of the proof of Theorem 2, eventually (almost) only zooming-in steps happen, and the rate of convergence is (no better than) k_{in} . The choice of k_{out} plays essentially no role. If instead $k_{\text{in}} < \rho$, the state does not contract fast enough for a stationary sequence of zooming-ins to establish. Then, zooming-in and zooming-out steps alternate in a complicated way, leading anyhow to consensus, with an evolution which depends on both k_{in} and k_{out} : the convergence can be faster than in the ideal case. Such phenomenon is rather surprising at a first glance, but we can explain it intuitively if we consider that the zooming-in/zooming-out algorithm takes advantage of the use of memory, and thus the sent messages are actually more informative.

Most of the simulations regard ring graphs, because concentrating on a simple example of a poorly connected graph helps to highlight the features of the algorithm. Weights of matrices P have been chosen according to a maximum degree rule. Hence a ring of $N = 20$ nodes induces a matrix whose essential spectral radius is $\rho = 0.9673$. The initial conditions have been generated according to a gaussian distribution, and all simulations with 20 nodes start from the same initial condition.

The transition between the two regimes is explored in Figure 1, confirming the role of ρ . Instead, Figures 2 and 2 consider the role of the number of quantization levels m . In the $k_{\text{in}} > \rho$ regime, at least $m = 3$ is required for convergence, and increasing m leads to improve the rate, approaching k_{in} . In the $k_{\text{in}} < \rho$ regime, $m = 1$ is enough for convergence, and increasing m leads to slightly improve the rate. Since $|\mathcal{S}_m| = m+2$, then 5 and 3 symbols are sufficient for convergence, respectively. If, taking advantage of the synchrony, one encodes the middle level, which is zero whenever m is odd, as a no-signal, the required number of level can be obtained with only two bits (one bit, respectively).

The scaling properties of the algorithm are a main point of interest. Figure 4 considers how the performance of the algorithm (with fixed parameters) degrades as N increases, in a sequence of ring graphs. The plot shows that m need not to depend on N as required Theorem 2. Instead, Figure 5 considers two cases in which the zooming factors k_{in} and k_{out} are depends on the graph, in such a way to keep the evolution in the same regime, in spite of the growth

of N .




Figure 1: Ring graph. Dependence on k_{in} : remark the threshold $\rho = 0.9673$. $N = 20$, $m = 6$, $k_{\text{out}} = 2$.

Not to restrict ourselves to ring graphs, we show some simulations regarding random geometric graphs in Figure 6. The plots show how the performance depends on the topology of the geometric graph: different samples have different spectral radiuses, and then different convergence rates in the ideal case. Instead, in most cases the zooming algorithm is in the $k_{\text{in}} < \rho$ regime, and then the performance is roughly independent of the sample. This suggests that in this regime the algorithm is little sensitive to the spectral radius, and rises the question of which is a significant graph theoretic index.

6 Conclusions

With theoretical and experimental results we showed that using the zooming scheme the average consensus problem can be efficiently solved although the agents can send only quantized information. Indeed the systems converge to average consensus at exponential speed, with a convergence rate which can be chosen as small as the convergence rate of consensus with perfect exchange of information, provided the number of quantization levels is large enough. Though the theoretical results are quite conservative the efficiency of these methods is apparent from simulations, which show the method to converge outside the scope of current results, and converge at a rate faster than the ideal algorithm. Improving the theoretical analysis is thus a major concern, especially about the issue of scalability in the number of agents.




Figure 2: Ring graph, in the $k_{\text{in}} > \rho$ regime. $N = 20$, $k_{\text{in}} = 0.97$, $k_{\text{out}} = 2$.

We have already underlined in Remark 1 that the algorithm requires the states estimates to be shared among neighbors. This implies that the algorithm is not suitable, in the present form, for implementation on a digital noisy channel, unless we allow a feedback on the channel state. A natural development of this research is thus designing algorithms to solve the average consensus problem at exponential speed over a network of digital noisy channels.

References

- [1] T. Aysal, M. Coates, and M. Rabbat. Distributed average consensus with dithered quantization. *IEEE Transactions on Signal Processing*, 56(10):4905–4918, 2008.
- [2] G. Baldan and S. Zampieri. An efficient quantization algorithm for solving average-consensus problems. In *European Control Conference*, Budapest, Hungary, Aug. 2009. Submitted.
- [3] R. W. Brockett and D. Liberzon. Quantized feedback stabilization of linear systems. *IEEE Transactions on Automatic Control*, 45(7):1279–1289, 2000.
- [4] R. Carli, F. Bullo, and S. Zampieri. Quantized average consensus via dynamic coding/decoding schemes. *International Journal of Robust and Non-linear Control*, 2009. To appear.




Figure 3: Ring graph, in the $k_{\text{in}} < \rho$ regime. $N = 20$, $k_{\text{in}} = 0.9$, $k_{\text{out}} = 2$.

- [5] R. Carli, F. Fagnani, P. Frasca, T. Taylor, and S. Zampieri. Average consensus on networks with transmission noise or quantization. In *European Control Conference*, pages 1852–1857, Kos, Greece, June 2007.
- [6] R. Carli, F. Fagnani, P. Frasca, and S. Zampieri. Efficient quantization techniques in consensus algorithms. In *NeCST Workshop*, Nancy, France, June 2007.
- [7] R. Carli, F. Fagnani, P. Frasca, and S. Zampieri. Gossip consensus algorithms with quantized communication. *Automatica*, 2009. Submitted.
- [8] R. Carli, F. Fagnani, A. Speranzon, and S. Zampieri. Communication constraints in the average consensus problem. *Automatica*, 44(3):671–684, 2008.
- [9] R. Elsässer, B. Monien, and S. Schamberger. Distributing unit size workload packages in heterogeneous networks. *J. Graph Algorithms Appl.*, 10(1):51–68, 2006.
- [10] P. Frasca. *Distributed optimization with communication constraints*. PhD thesis, Politecnico di Torino, Mar. 2009.
- [11] P. Frasca, R. Carli, F. Fagnani, and S. Zampieri. Average consensus on networks with quantized communication. *International Journal of Robust and Nonlinear Control*, 2009. In press.
- [12] S. Kar and J. M. F. Moura. Distributed consensus algorithms in sensor networks: Quantized data. Dec. 2007. Submitted.




Figure 4: Ring graph, with $m = 3$, $k_{\text{in}} = 0.99$, $k_{\text{out}} = 2$. For $N > 20$ we are out of the scope of Theorem 2, and for $N > 30$, we have $k_{\text{in}} < \rho$.

- [13] A. Kashyap, T. Başar, and R. Srikant. Quantized consensus. *Automatica*, 43(7):1192–1203, 2007.
- [14] G. Nair, F. Fagnani, S. Zampieri, and R. Evans. Feedback control under data rate constraints: An overview. *Proceedings of the IEEE*, 95(1):108–137, Jan. 2007.
- [15] A. Nedic, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis. On distributed averaging algorithms and quantization effects. Technical Report 2778, Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Nov. 2007.
- [16] R. Olfati-Saber, J. A. Fax, and R. M. Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.
- [17] S. Tatikonda and S. Mitter. Control under communication constraints. *IEEE Transactions on Automatic Control*, 49(7):1056–1068, 2004.
- [18] L. Xiao, S. Boyd, and S. Lall. A scheme for robust distributed sensor fusion based on average consensus. In *Symposium on Information Processing of Sensor Networks (IPSN)*, pages 63–70, Los Angeles, CA, Apr. 2005.

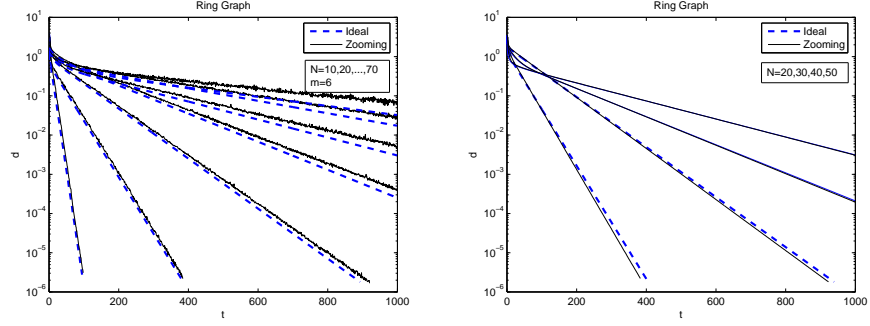


Figure 5: Ring graph, $m = 6$, for increasing N . Left plot assumes $k_{in} = 1.01\rho$ and $k_{out} = 2$, yielding the $k_{in} > \rho$ regime. Right plot assumes $k_{in} = 0.5\rho$, $k_{out} = k_{in}^{-1}$, yielding the $k_{in} < \rho$ regime.




Figure 6: Different performance on different samples of random geometric graph with $R = 0.5$. $m = 3$, $k_{in} = 0.3$, $k_{out} = 2$.